

**Gates & Cooper** LLPHoward Hughes Center  
6701 Center Drive West, Suite 1050  
Los Angeles, California 90045**RECEIVED  
CENTRAL FAX CENTER**

SEP 26 2007

**FAX TRANSMISSION TO USPTO**TO: Commissioner for Patents  
Attn: Examiner Cindy Nguyen  
Patent Examining Corps  
Facsimile Center  
Alexandria, VA 22313-1450FROM: George H. Gates  
OUR REF.: STL919990184US3  
TELEPHONE: (310) 642-4146Total pages, including cover letter: 11PTO FAX NUMBER: 571-273-8300


If you do NOT receive all of the pages, please telephone us at (310) 641-8797, or fax us at (310) 641-8798.

|                                |   |
|--------------------------------|---|
| Title of Document Transmitted: | TRANSMITTALS AND REPLY BRIEF OF APPELLANT   |
| Applicant:                     | David E. Simmen   |
| Serial No.:                    | 10/807,871  |
| Filed:                         | March 24, 2004  |
| Group Art Unit:                | 2161  |
| Title:                         | QUERY OPTIMIZATION TECHNIQUE FOR OBTAINING IMPROVED CARDINALITY ESTIMATES USING STATISTICS ON PRE-DEFINED QUERIES |
| Our Ref. No.:                  | STL919990184US3   |

Please charge all fees to Deposit Account No. 09-0460 of IBM Corporation, the assignee of the present application.

By: Name: George H. Gates  
Reg. No.: 33,500

I hereby certify that this paper is being transmitted by facsimile to the U.S. Patent and Trademark Office on the date shown below.

  
SignatureSeptember 26, 2007  
Date

GHG/cf

G&amp;C 30571.272-US-C1

**RECEIVED  
CENTRAL FAX CENTER**

SEP 26 2007

Duc Date: September 26, 2007

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

|             |  |                 |                 |
|-------------|--|-----------------|-----------------|
| Applicant:  | David E. Simmen  | Examiner:       | Cindy Nguyen    |
| Serial No.: | 10/807,871   | Group Art Unit: | 2161            |
| Filed:      | March 24, 2004   | Docket:         | STL919990184US3 |
| Title:      | QUERY OPTIMIZATION TECHNIQUE FOR OBTAINING IMPROVED CARDINALITY<br>ESTIMATES USING STATISTICS ON PRE-DEFINED QUERIES |                 |                 |

**CERTIFICATE OF MAILING OR TRANSMISSION UNDER 37 CFR 1.8**

I hereby certify that this correspondence is being filed via facsimile transmission to the U.S. Patent and Trademark Office on September 26, 2007.

By: \_\_\_\_\_  
Name: Christine Flores

Commissioner for Patents  
P.O. Box 1450  
Alexandria, VA 22313-1450

Dear Sir:

We are transmitting herewith the attached:

- ☒ Transmittal sheet, in duplicate, containing a Certificate of Mailing or Transmission under 37 CFR 1.8.
- ☒ Reply Brief of Appellant(s).

Please consider this a **PETITION FOR EXTENSION OF TIME** for a sufficient number of months to enter these papers, if appropriate.

Please charge all fees to Deposit Account No. 09-0460 of IBM Corporation, the assignee of the present application. A duplicate of this paper is enclosed.

Customer Number 45729**GATES & COOPER LLP**

Howard Hughes Center  
6701 Center Drive West, Suite 1050  
Los Angeles, CA 90045  
(310) 641-8797

By: George H. Gates  
Name: George H. Gates  
Reg. No.: 33,500  
GHG/cf

**RECEIVED  
CENTRAL FAX CENTER****SEP 26 2007**

Due Date: September 26, 2007

**IN THE UNITED STATES PATENT AND TRADEMARK OFFICE**

|             |   |                 |                 |
|-------------|---|-----------------|-----------------|
| Applicant:  | David E. Simmen   | Examiner:       | Cindy Nguyen    |
| Serial No.: | 10/807,871  | Group Art Unit: | 2161            |
| Filed:      | March 24, 2004  | Docket:         | STL919990184US3 |
| Title:      | QUERY OPTIMIZATION TECHNIQUE FOR OBTAINING IMPROVED CARDINALITY ESTIMATES USING STATISTICS ON PRE-DEFINED QUERIES |                 |                 |

**CERTIFICATE OF MAILING OR TRANSMISSION UNDER 37 CFR 1.8**

I hereby certify that this correspondence is being filed via facsimile transmission to the U.S. Patent and Trademark Office on September 26, 2007.

By: \_\_\_\_\_  
Name: Christine Flores

Commissioner for Patents  
P.O. Box 1450  
Alexandria, VA 22313-1450

Dear Sir:

We are transmitting herewith the attached:

- ☒ Transmittal sheet, in duplicate, containing a Certificate of Mailing or Transmission under 37 CFR 1.8.
- ☒ Reply Brief of Appellant(s).

Please consider this a **PETITION FOR EXTENSION OF TIME** for a sufficient number of months to enter these papers, if appropriate.

Please charge all fees to Deposit Account No. 09-0460 of IBM Corporation, the assignee of the present application. A duplicate of this paper is enclosed.

Customer Number 45729  
GATES & COOPER LLP  
Howard Hughes Center  
6701 Center Drive West, Suite 1050  
Los Angeles, CA 90045  
(310) 641-8797

By: GHG  
Name: George H. Gates  
Reg. No.: 33,500  
GHG/cf

**RECEIVED  
CENTRAL FAX CENTER**

SEP 26 2007

Due Date: September 26, 2007

**IN THE UNITED STATES PATENT AND TRADEMARK OFFICE  
BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES**

|                                     |   |                        |
|-------------------------------------|---|------------------------|
| In re Application of:               | ) |                        |
|                                     | ) |                        |
| Inventor: David E. Simmen           | ) | Examiner: Cindy Nguyen |
|                                     | ) |                        |
| Serial #: 10/807,871                | ) | Group Art Unit: 2161   |
|                                     | ) |                        |
| Filed: March 24, 2004               | ) | Appeal No.: _____      |
|                                     | ) |                        |
| Title: QUERY OPTIMIZATION TECHNIQUE | ) |                        |
| FOR OBTAINING IMPROVED              | ) |                        |
| CARDINALITY ESTIMATES USING         | ) |                        |
| STATISTICS ON PRE-DEFINED           | ) |                        |
| QUERIES                             | ) |                        |

**REPLY BRIEF OF APPELLANT****MAIL STOP APPEAL BRIEF - PATENTS**

Commissioner for Patents  
P.O. Box 1450  
Alexandria, VA 22313-1450

Dear Sir:

In accordance with 37 C.F.R. §41.41, Appellant's attorney hereby submits the Reply Brief of Appellant in response to the Examiner's Answer dated July 26, 2007 for the above-identified application.

No fee is required for filing this Reply Brief. However, the Office is authorized to charge any necessary fees or credit any overpayments to Deposit Account No. 09-0460 of IBM Corporation, the assignee of the present application.

**I. ARGUMENTS**

In the Answer, the Examiner essentially reiterates the prior rejections, albeit with additional remarks. In this regard, this Reply Brief of Appellant incorporates by reference herein the entirety of the previously filed Brief of Appellant. Moreover, additional arguments are also presented below.

A. Arguments directed to the first grounds for rejection: Whether claims 1, 2-5, 12-17, and 23-28 are anticipated under 35 U.S.C. §102(e) by U.S. Publication No. 2003/0088558 (Zaharioudakis).

1. Claims 1, 12 and 23

With regard to independent claims 1, 12 and 23, the Examiner's Answer restates the rejections of the final Office Action. Consequently, the Reply Brief of Appellant incorporates by reference herein the entirety of the arguments in the previously filed Brief of Appellant.

2. Claims 2, 13 and 24

With regard to dependent claims 2, 13, and 24, which recite that each of the pre-defined queries is associated with an automatic summary table, a materialized view or a view, these claims stand or fall with independent claims 1, 12 and 23.

3. Claims 3, 14 and 25

With regard to dependent claims 3, 14, and 25, which recite generating cardinality estimates for one or more query execution plans for the query using the statistics of one or more of the pre-defined queries that vertically overlap the query, and using the generated cardinality estimates to determine an optimal query execution plan for the query, the Examiner's Answer restates the rejections of the final Office Action. Consequently, the Reply Brief of Appellant incorporates by reference herein the entirety of the arguments in the previously filed Brief of Appellant.

4. Claims 4, 15 and 26

With regard to dependent claims 4, 15, and 26, which recite that the statistics are used to improve a combined selectivity estimate of one or more predicates of the query, the Examiner's Answer restates the rejections of the final Office Action. Consequently, the Reply Brief of Appellant incorporates by reference herein the entirety of the arguments in the previously filed Brief of Appellant.

5. Claims 5, 16 and 27

With regard to dependent claims 5, 16, and 27, which recite that the predicates are applied by one or more of the pre-defined queries, these claims stand or fall with independent claims 1, 12 and 23.

B. Arguments directed to the second grounds for rejection: Whether claims 6-11, 17-22 and 28-33 are obvious under 35 U.S.C. §103(a) over U.S. Publication No. 2003/0088558 (Zaharioudakis) in view of U.S. Patent No. 6,496,819 (Bello).

1. Claims 6, 17 and 28

With regard to dependent claims 6, 17 and 28, which recite that the selectivity estimate comprises a ratio of a cardinality of the pre-defined query to a product of cardinalities of base tables referenced in the pre-defined query and the query, the Examiner's Answer restates the rejections of the final Office Action, but also includes a citation to a new location in Bello, namely col. 10, lines 11-12, as well as the originally cited location in Bello, namely col. 10, lines 45-67. These portions of Bello are reproduced below:

Bello: col. 10, lines 11-12 (actually, lines 8-18)

Various criteria may be used during this pruning process. For example, one possible pruning criteria may be that at least one of the tables referenced in the received query must be a base table of the materialized view. Based on this criteria, a materialized view that has base tables A, B, and C would qualify as a "possible materialized view" with respect to a query that requires a join between tables A and D. On the other hand, a materialized view that has base tables B, C and E would not qualify as a "possible materialized view" with respect to a query that requires a join between tables A and D.

Bello: col. 10, lines 45-67

According to one embodiment of the invention, the query reduction factor for a materialized view is the ratio of (1) the sum of the cardinalities of matching relations in the query that will be replaced by the materialized view to (2) the cardinality of the materialized view. When use of the materialized view requires a join back, then the query reduction factor is adjusted to take into account the overhead associated with the join back. This is simply one example of how the query reduction factor may be determined. The present invention is not limited to any particular query reduction factor calculation.

At step 216, the query reduction factor for the materialized view being processed is compared to the highest previously calculated query reduction factor (the query reduction factor of the "current best materialized view"). If the query reduction factor for the materialized view being processed is greater than the highest previously calculated query reduction factor, then the materialized view being processed becomes the new "current best materialized view" (step 218) and control passes back to step 202. Otherwise, the materialized view being processed is removed from consideration (step 230) and control passes back to step 202.

The above portions of Bello merely disclose the use of pruning criteria for determining the eligibility of materialized views, as well as a query reduction factor, which is the ratio of (1) the sum

of the cardinalities of matching relations in the query that will be replaced by the materialized view to (2) the cardinality of the materialized view.

However, the query reduction factor is not the same as the selectivity estimate, which comprises a ratio of a cardinality of the pre-defined query to a product of cardinalities of base tables referenced in the pre-defined query and the query. Instead, these are completely different values.

## 2. Claims 7, 18 and 29

With regard to dependent claims 7, 18 and 29, which recite that zero or more predicates of the query are applied by one of the pre-defined queries and wherein the remaining predicates are eligible to be applied on the pre-defined query, the Examiner's Answer restates the rejections of the final Office Action, but also includes a citation to a new location in Bello, namely col. 8, lines 62-67, as well as the originally cited location in Bello, namely col. 10, lines 20-45. These portions of Bello are reproduced below:

### Bello: col. 8, lines 62-67

For example, assume that a query contains the join  $A > < B > < C -> D$  and a materialized view definition contains the join  $A > < B > < E$ . Under these conditions, the join  $A > < B$  is the common section, the joins  $B > < C$  and  $C -> D$  belong to the query delta, and the join  $B > < E$  belongs to the materialized view delta.

### Bello: col. 10, lines 20-45

Steps 202 and 204 form a loop in which each materialized view in the set of possible materialized views is processed. During the processing of each materialized view, the database server determines whether the materialized view is actually eligible to be used in a rewrite of the received query, and if so, the relative benefit gained by using that particular materialized view. Specifically, at step 202, it is determined whether any "possible materialized views" are left to be processed. If all possible materialized views have been processed, control passes to step 220. Otherwise, control passes to step 204.

At step 204, an unprocessed possible materialized view is selected. At step 206, it is determined whether the selected materialized view is eligible for use in rewriting the received query. If the selected materialized view is not found to be eligible, control passes to step 230, and the materialized view is removed from consideration. From step 230, control passes back to step 202.

If the selected materialized view is found to be eligible at step 206, control passes to step 214. At step 214, a "query reduction factor" is computed for the materialized view currently being processed. The query reduction factor is a value that estimates how useful it will be to access the materialized view to process the received query. The higher the query reduction factor, the greater the estimated benefit of using the materialized view to process the query.

The above portions of Bello merely disclose that a query and materialized view may have a common section as well as deltas (differences), and that a database server determines whether a materialized view is eligible to be used in a rewrite of a query.

However, this determination does not include a determination whether zero or more predicates of the query are applied by one of the pre-defined queries and wherein the remaining predicates are eligible to be applied on the pre-defined query. Instead, these are completely different determinations.

3. Claims 8, 19 and 30

With regard to dependent claims 8, 19 and 30, which recite that a predicate is eligible to be applied on the pre-defined query if it can be evaluated using the output columns and expressions of the pre-defined query, these claims stand or fall with claims 7, 18 and 29.

4. Claims 9, 20 and 31

With regard to dependent claims 9, 20 and 31, which recite determining a subpredicate combined selectivity estimate of the unapplied eligible predicates using column distribution statistics of the pre-defined query, the Examiner's Answer omits the rejections of the final Office Action, but uses citations to new locations in Bello, namely col. 19, lines 40-55, as compared to the originally cited location in Bello, namely col. 10, lines 30-36. These portions of Bello are reproduced below:

Bello: col. 19, lines 40-55 (actually, lines 36-54)

Because rows from fact 900 may be produced more than once during the join, they may be counted more than once during the aggregation. Hence, the resulting aggregate values are not accurate. For example, the actual sum of sales for stores 1, 2 and 3 are \$12, \$18 and \$12, respectively. However, the values in the sumsales column of summary table 904 for stores 1, 2, and 3 are \$24, 54, and \$12, respectively.

FIG. 9C illustrates a summary table 906 that is generated by the summary table definition (D2):

```
ST (store, sumsales, scalingfactor, city):  
select fact.store, sum(fact.sales),  
count(distinct decode(detail.rowid, null, '1', fact.store .parallel. detail.rowid))  
detail.city  
from fact, detail  
where fact.store=detail.store (+)  
group by fact.store, detail.city;
```



The above portions of Bello merely disclose that aggregate values may be incorrect because rows are counted more than once during a join.

However, nothing in the cited portions of Bell teach or suggest determining a subpredicate combined selectivity estimate of the unapplied eligible predicates using column distribution statistics of the pre-defined query.

5. Claims 10, 21 and 32

With regard to dependent claims 10, 21 and 32, which recite that a cardinality ratio comprises a ratio of a cardinality of the pre-defined query to a product of cardinalities of base tables referenced in the pre-defined query and the query, the Examiner's Answer restates (and also changes) the rejections of the final Office Action, but also includes a citation to a new location in Bello, namely col. 10, lines 11-12, as well as the originally cited location in Bello, namely col. 10, lines 45-67 (originally cited as col. 10, lines 37-56). These portions of Bello are reproduced below:

Bello: col. 10, lines 11-12 (actually, lines 8-18)

Various criteria may be used during this pruning process. For example, one possible pruning criteria may be that at least one of the tables referenced in the received query must be a base table of the materialized view. Based on this criteria, a materialized view that has base tables A, B, and C would qualify as a "possible materialized view" with respect to a query that requires a join between tables A and D. On the other hand, a materialized view that has base tables B, C and E would not qualify as a "possible materialized view" with respect to a query that requires a join between tables A and D.

Bello: col. 10, lines 45-67

According to one embodiment of the invention, the query reduction factor for a materialized view is the ratio of (1) the sum of the cardinalities of matching relations in the query that will be replaced by the materialized view to (2) the cardinality of the materialized view. When use of the materialized view requires a join back, then the query reduction factor is adjusted to take into account the overhead associated with the join back. This is simply one example of how the query reduction factor may be determined. The present invention is not limited to any particular query reduction factor calculation.

At step 216, the query reduction factor for the materialized view being processed is compared to the highest previously calculated query reduction factor (the query reduction factor of the "current best materialized view"). If the query reduction factor for the materialized view being processed is greater than the highest previously calculated query reduction factor, then the materialized view being processed becomes the new "current best materialized view" (step 218) and control passes back to step 202. Otherwise, the materialized view being processed is removed from consideration (step 230) and control passes back to step 202.

The above portions of Bello merely disclose the use of pruning criteria for determining the eligibility of materialized views, as well as a query reduction factor, which is the ratio of (1) the sum of the cardinalities of matching relations in the query that will be replaced by the materialized view to (2) the cardinality of the materialized view.

However, the query reduction factor is not the same as the cardinality ratio, which comprises a ratio of a cardinality of the pre-defined query to a product of cardinalities of base tables referenced in the pre-defined query and the query.

6. Claims 11, 22 and 33

With regard to dependent claims 11, 22 and 33, which recite that the selectivity estimate comprises a product of the subpredicate combined selectivity estimate and the cardinality ratio, the Examiner's Answer omits the rejections of the final Office Action, but uses citations to new locations in Bello, namely col. 10, lines 45-67 and col. 10, lines 11-12, as compared to the originally cited location in Bello, namely col. 11, line 55 to col. 12, line 41. These portions of Bello are reproduced below:

Bello: col. 10, lines 11-12 (actually, lines 8-18)

Various criteria may be used during this pruning process. For example, one possible pruning criteria may be that at least one of the tables referenced in the received query must be a base table of the materialized view. Based on this criteria, a materialized view that has base tables A, B, and C would qualify as a "possible materialized view" with respect to a query that requires a join between tables A and D. On the other hand, a materialized view that has base tables B, C and E would not qualify as a "possible materialized view" with respect to a query that requires a join between tables A and D.

Bello: col. 10, lines 45-67

According to one embodiment of the invention, the query reduction factor for a materialized view is the ratio of (1) the sum of the cardinalities of matching relations in the query that will be replaced by the materialized view to (2) the cardinality of the materialized view. When use of the materialized view requires a join back, then the query reduction factor is adjusted to take into account the overhead associated with the join back. This is simply one example of how the query reduction factor may be determined. The present invention is not limited to any particular query reduction factor calculation.

At step 216, the query reduction factor for the materialized view being processed is compared to the highest previously calculated query reduction factor (the query reduction factor of the "current best materialized view"). If the query reduction factor for the materialized view being processed is greater than the highest

previously calculated query reduction factor, then the materialized view being processed becomes the new "current best materialized view" (step 218) and control passes back to step 202. Otherwise, the materialized view being processed is removed from consideration (step 230) and control passes back to step 202.

The above portions of Bello merely disclose the use of pruning criteria for determining the eligibility of materialized views, as well as a query reduction factor, which is the ratio of (1) the sum of the cardinalities of matching relations in the query that will be replaced by the materialized view to (2) the cardinality of the materialized view.

However, the query reduction factor is not the same as the selectivity estimate, which comprises a product of the subpredicate combined selectivity estimate and the cardinality ratio.

### III. CONCLUSION

In light of the above arguments, Appellant's attorney respectfully submits that the cited references do not anticipate nor render obvious the claimed invention. More specifically, Appellant's claims recite novel physical features which patentably distinguish over any and all references under 35 U.S.C. §§ 102 and 103.

As a result, a decision by the Board of Patent Appeals and Interferences reversing the Examiner and directing allowance of the pending claims in the subject application is respectfully solicited.

Respectfully submitted,

GATES & COOPER LLP  
Attorneys for Appellant

Howard Hughes Center  
6701 Center Drive West, Suite 1050  
Los Angeles, California 90045  
(310) 641-8797

Date: September 26, 2007

By: George H. Gates  
Name: George H. Gates  
Reg. No.: 33,500

GHG/